

UNITED STATES PATENT APPLICATION
FOR
METHOD AND APPARATUS FOR PROVIDING DYNAMICAL
SCALABILITY

INVENTORS:

TONY HAMILTON

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(408) 720-8598

"Express Mail" mailing label number EL627466588US

Date of Deposit: September 30, 2000

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231

Michelle Begay

(Typed or printed name of person mailing paper or fee)

Michelle Begay

(Signature of person mailing paper or fee)

September 30, 2000

Date

[illegible]

This invention relates to computers. Specifically, this invention relates to providing dynamical scalability to portable computers, including notebook computers.

Portable systems, such as notebook computers, have steadily become more powerful with improved technology. Notebook computers are popular because they are mobile and can be used in remote places such as on airplanes. They can also be used in homes and offices with accesses to AC outlets.

As the processor and bus speeds steadily rise and the graphics performances steadily improve, the power dissipation of the system components to support them also increases. Power dissipation is generally not an issue when the notebook computer is connected to an external power source such as an AC outlet because the power supply is infinite. When the notebook computer relies on a battery as its power source, however, power dissipation becomes important because the battery has a finite life span after which it needs to be recharged.

1

reasonable amount of time to use his notebook computer in a remote place such as on an airplane.

One solution to maximize battery life is offered by the Geyserville™ processor technology developed by the Intel Corporation located in Santa Clara, California. Geyserville extends the battery life by providing a dynamically scalable processor performance. Dynamic scaling refers to automatically adjusting the processor's performance state (processor's clock frequency and supply voltage) depending on the power source. Power dissipation of the processor is proportional to its clock frequency and to the square of its supply voltage. By slowing down the processor's clock and reducing the its supply voltage when the portable system is relying on the battery for its power source, the Geyserville technology dramatically reduces the processor's power dissipation.

The Geyserville solution is only a partial one however, and there is a further need to increase battery life span.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

- 5 Figure 1 is a block diagram of a system in which the method according to one embodiment is implemented.

Figure 2 is a flow diagram of the method according to one embodiment.

DETAILED DESCRIPTION

The present invention extends the dynamic scalability of the Geyserville technology beyond the processor. In one embodiment, system buses and other system components are dynamically scaled.

5 Scalability refers to increasing or decreasing certain attributes of a system's components such as clock frequencies, supply voltages, graphics performances, buffer strengths, and others. Dynamic scalability refers to scalability upon occurrence of a predetermined event, such as a power management event, without user intervention. In the ensuing description,
10 each attribute is described in relation to two performance states (high level and low level) between which scaling is performed. It is contemplated, however, that each attribute can be scaled among multiple performance states.

Fig. 1 is a block diagram of a system 10 in which the method
15 according to one embodiment of the present invention is implemented. In one embodiment, the system 10 is a computer. In other embodiments, system 10 can be a portable computer, a hand-held electronic device, and the like. In one embodiment, the computer 10 includes a System Management Controller (SMC) 12. The SMC 12 controls the transfer of data between the
20 chipsets and the peripheral devices such as the processor 16, the Input Output Control Hub (I/O CH) 30, the Graphics Memory Control Hub (GMCH) 26, the keyboard 34, and the Input Output Device (I/O Device) 36.

In one embodiment, the SMC 12 is formed of one or more layers including, for example, software, firmware, and hardware. In one
25 embodiment, the SMC 12 is connected to the chipsets and the peripheral devices by a first bus 14. In one embodiment, the first bus is referred to as the System Management Bus (SMBUS). In one embodiment, the chipsets include the GMCH 26 and the I/O CH 30. In one embodiment, the chipsets

provide the electronic interfaces between the processor 16, the memory sub system 20, and the graphics sub system 42.

A bus is a collection of wires through which data is transferred from one part of a computer to another. Two characteristics of every bus are its
5 width and its clock speed. The width determines how much data can be transmitted at one time. For example, a 16-bit bus can transfer 16 bits of data, whereas a 32-bit bus can transmit 32 bits of data. The clock speed is measured in megahertz (MHZ) and indicates how fast the data is transferred. A higher bus clock speed requires a higher minimum bus
10 supply voltage. If the data transmission load on the SMBUS 14 is low, for example, because of low activity, lowering the clock speed and the supply voltage of the SMBUS 14 can reduce the power dissipation of the SMBUS 14.

As stated previously, reducing the power dissipation is generally not an issue when the computer 10 is connected to an external power source
15 such as an AC outlet because the power supply is infinite. When the computer 10 is connected to a battery, however, reducing the power dissipation provides a benefit in the form of increased battery longevity.

In one embodiment, the computer 10 includes an additional bus such as the Memory Bus (MBUS) 34. In one embodiment, the MBUS 34 connects
20 the memory sub system 20 to the GMCH 26. In one embodiment, the MBUS 34 operates at a higher speed than the SMBUS 14.

Fig. 2 is a flow diagram of the method according to one embodiment to be implemented in the computer 10 of Fig. 1.

In one embodiment, the SMC 12 detects the predetermined power
25 management event (at 102). In one embodiment, the power management event includes a change in the power source of the computer 10 from an external source to a battery or vice versa. Upon detecting the power management event, the SMC 12 automatically places the computer in a Low Activity State (at 104). In one embodiment, the Low Activity State includes

the Deep Sleep Power Management State (ACPI Specification C3 State).

ACPI Specification stands for Advanced Configuration and Power Interface Specification, Revision 2.0, published on July 27, 2000. During the C3 State, the clock generator 18 input to the components is disabled. That stops all logical functions processing inside the components but allows the internal settings to be maintained (e.g., values stored in registers and caches).

Following, the SMC 12 utilizes a management command structure and an algorithm to send multiplexed commands by utilizing the SMBUS 14 to perform various scaling operations (at 106, 108, 110 and 112). Sending multiplexed commands refers to combining the command signals into one signal for transmission to the multiple components. Two communications devices are used for multiplexing: a multiplexor residing in the SMC 12 to combine the command signals into one signal and a demultiplexor residing in the component unit to separate the command signal pertaining to that component. Sending commands in a multiplexed fashion allows for the scaling operations (at 106-112) to occur simultaneously. In an alternative embodiment, the SMC 12 sends individual commands instead of multiplexed commands.

In one embodiment, performing a scaling operation includes changing the setting of a performance register inside the component. The algorithm contains the predetermined performance register settings and is automatically executed by the SMC 12 upon the occurrence of the power management event. In one embodiment, the algorithm has two predetermined settings (high level and low level) for each of the attributes to be scaled. In other embodiments, programmable devices such as fuse banks and non-volatile memory are used instead of the performance registers.

In one embodiment, the frequencies of the processor clock, the memory sub system clock, and the clock generator 18 are scaled between a high level and a low level (at 106) depending on the power management

event. In one embodiment, the high level is 133 MHZ and the low level is 100 MHZ. The clock frequencies are set to the low level if the computer 10 power source was changed from the outside source to the battery. The clock speeds are set to the high level if the computer 10 power source was

5 changed from the battery to the outside source.

In one embodiment, the setting of the performance register represents bus fraction data specifying the ratio of the component's clock frequency to the clock generator 18 output.

In one embodiment, the buffer strength of the chipset 40 is scaled

10 between a high level and a low level (at 108). The chipset buffers drive the SMBUS 14 and the MBUS 34 by reading data from and writing data onto the buses 14 and 34. The chipset buffers act as holding areas, enabling the processor 16 to manipulate data before transferring it to the buses 14 and 34. The chipset buffers compensate for the speed differences between the

15 processor 16 and the buses 14 and 34. Buffer strength can be scaled between a high level and a low level, based on the data transmission load and clock frequencies of the buses 14 and 34. A higher buffer strength requires a higher minimum bus supply voltage. By scaling the buffer strength to the low level when the computer 10 is using the battery as its power source,

20 power dissipation is reduced. In one embodiment, the buses speeds are scaled (at 110) between a high level and a low level. In one embodiment, the high level is 133 MHZ and the low level is 100 MHZ. In another embodiment, the high level is 450 MHZ and the low level is 350 MHZ. The bus speeds of the SMBUS 14 and the MBUS 34 are scaled by scaling their

25 respective clock frequencies.

A minimum level of supply voltage is required to support the components clock frequencies and buffer strengths. In the embodiment, depending on the clock frequencies and buffer strengths, supply voltages of the respective components are scaled between a high level and a low level

(at 112). In one embodiment, the high level is 1.8 VDC and the low level is 1.3 VDC.

In another embodiment, the high level is 14.0 VDC and the low level is 8.0 VDC.

5 In one embodiment, the graphics performances are scaled between a high level and a low level. The high level includes the 4X performance level under the Accelerated Graphics Port (AGP) Interface Specification developed by the Intel Corporation. The low level includes the 2X performance level under the AGP Specification. The 4X performance level
10 requires higher minimum bus speeds and supply voltages than the 2X performance level. It is contemplated that in another embodiment, the low level includes the AGP Specification level 1X.

The present invention thus extends the scalability feature of Geyserville beyond the processor to various components across the computer 10 platform. By scaling to low levels various attributes of the components when a battery power source is utilized, power dissipation of the computer 10 is significantly reduced.

The SMC 12 next determines if all scaling operations have been successfully performed (at 114). In one embodiment, predefined registers including the model-specific registers (MSRs) or component ID registers are read to make that determination. The predefined registers may be updated by the components after the scaling operations have been successfully performed. If the scaling operations were successful, the SMC 12 removes the computer 10 from the Low Activity State (at 116).

25 In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set

forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

In addition, the methods as described above can be stored in memory of a computer system as a set of instructions to be executed. In addition, the instructions to perform the methods as described above could alternatively be stored on other forms of computer-readable mediums, including magnetic and optical disks. For example, the method of the present invention can be stored on computer-readable mediums, such as magnetic disks or optical disks that are accessible via a disk drive (or computer-readable medium drive).

063-003